arXiv:1402.5207v1 [cs.DC] 21 Feb 2014

NOISEOFF: A Backoff Protocol for a Dynamic, Noisy World

Michael A. Bender[1,2], Jeremy T. Fineman[3], Seth Gilbert[4], and Maxwell Young[5]

[1]Department of Computer Science, Stony Brook University, Stony Brook, NY 11794-4400, USA
bender@cs.stonybrook.edu

[2]Tokutek, Inc.,
http://tokutek.com

[3]Department of Computer Science, Georgetown University, USA
jfineman@cs.georgetown.edu

[4]Department of Computer Science, National University of Singapore, Singapore
seth.gilbert@comp.nus.edu.sg

[5]Department of Computer Science, Drexel University, USA
myoung@cs.drexel.edu

Abstract

This paper gives a dynamic backoff protocol, called NOISEOFF, that provides constant throughput with online packet arrivals. NOISEOFF tolerates disruptive jamming/noise, and by leveraging low-power monitoring, it ensures low energy usage.

The packet arrivals are determined adaptively by an adversary, who also incarnates all faulty devices in the system and uses them to launch attacks. The paper analyzes the cases where the total number of packets is (1) a finite number n and (2) infinite. In the first case, NOISEOFF achieves constant expected throughput and, in the second, constant throughput with probability 1.

NOISEOFF is provably robust to adaptive adversarial jamming. To model the energy constraints of these networks, there is an associated cost of 1 to access the channel in a time slot, and a cost of 0 for low-power monitoring that can determine if the channel is busy (but does not receive messages). For an adversary who incurs a cost of J to jam, NOISEOFF provides the following guarantees. When the number of packets is a finite n, the average expected cost for successfully sending a packet is $O(\log^2(n+J))$. In the infinite case, the average expected cost for successfully sending a packet is $O(\log^2(\eta) + \log^2(J))$ where η is the maximum number of packets that are ever in the system concurrently.

This research was supported in part by NSF grants CCF 1114809, CCF 1217708, CCF 1218188, IIS 1247726, and IIS 1251137.

1 Introduction

Randomized backoff is a powerful approach for coordinating access to a shared resource. In this approach, processes repeatedly attempt to access the resource. If two processes collide while vying for the resource, then the access fails, and the processes independely wait for random amounts of time before retrying. After each collision, a process's waiting time increases, resulting in reduced contention for the resource and a greater probability of success.

In wireless networks, randomized backoff is commonly employed by medium access control (MAC) protocols as a mechanism for reducing contention on a shared wireless communication channel. Backoff protocols arise in most areas of networking (e.g., LANs, WANs, sensor networks, wireless networks), as well as in transactional memory [35], and speculative-lock elision [49]. Given the range of backoff applications, it should not be surprising that there are many flavors and properties of backoff.

This paper focuses on three desirable properties, motivated by the wireless setting:

- **Dynamic (Online) Packet Scheduling.** The network traffic may fluctuate as the set of demands from each user change. For example, crowds of users with hand-held devices may pass through an area of service, or wireless-enabled vehicles may congregate near a base station in a city road network. Changes in network traffic can congest the shared medium and negatively impact throughput.
- **Noise and Jamming resistance.** Scheduling access can be problematic because of interference from faulty devices. These faults may arise as a result of hardware failures, software bugs, or even malicious behavior. From an algorithmic perspective, we can model faulty devices by assuming a single ***adversary*** who coordinates their actions to cause worst-case interference. Such an attack is referred to as ***jamming***, and it poses a well-known threat to the availability of wireless networks [58, 60].
- **Energy efficiency.** Devices are typically battery powered, and replenishing this onboard supply may be difficult, or even impossible, depending on the deployment. A useful measure of energy usage is the number of slots during which a device sends or receives packets. For example, the CC2420 radio on the popular Telos mote has sending and receiving costs of 35mW (at 0 dBm) and 38mW, respectively, while the energy-efficient sleep and idle states use on the order of μW [46].[1]

While there is a deep and impressive algorithmic literature on achieving some of these properties, there are few (if any) provably efficient protocols that guarantee all three properties.

There is a long history of backoff protocols used in practice, with binary exponential backoff the most common.[2] For a long time, most of the analytic results on backoff protocols assumed statistical queueing theory models and focused on the question of what packet-arrival rates are stable (see [26, 27, 29, 33, 34, 48]).

More recently, there has been work on "adversarial queueing theory," looking at the worst-case performance of these protocols [2, 7, 8, 13, 14, 20, 25, 28, 30, 31, 59]. A common theme throughout these papers, however, is that dynamic arrivals are hard to cope with. When all the packets begin at the same time, very efficient protocols are possible. When packets begin at different times, the problem is much harder. This has been explicitly studied in the context of the "wake-up problem" [11, 12, 15], which looks at how long it takes for a single transmission to succeed when packets arrive dynamically. In contrast, we focus here on guaranteeing constant throughput with dynamic worst-case packet arrivals.

As the focus of research on backoff protocols has shifted to wireless networks, there has been an increasing emphasis on coping with noise and jamming: can we still reduce contention effectively in a noisy environment? In a surprising breakthrough, Awerbuch et al. [6] showed that good throughput is indeed possible, even if noise or an attacker causes disruption for a constant fraction of the execution. Their protocol assumes a fixed number of "stations" that are continually transmitting packets (i.e., arrivals are not

[1] Similar relative costs hold for the older Mica motes [16, 17]. Energy seems likely to remain a scarce resource [3].

[2] Interestingly, binary exponential backoff [41], the most famous backoff protocol, does not even have constant throughput for batch problems, and with bursty/adversarial packet arrivals it performs even worse [7]. Even with Poisson arrivals, there are better protocols, such as polynomial backoff [33].

dynamic). A number of elegant results have followed [44, 51–55]. At the heart of these approaches is an attempt to maintain a good estimate of the number of stations, and hence the proper probability with which to broadcast, via a multiplicative-weights technique. (As a result, even though they assume a fixed number of stations, their protocols can adapt to reasonable changes as long as they do not happen too fast.) Our goal here is to tolerate jamming and noise, while allowing for fully dynamic packet arrivals.

Energy has also become increasingly important, and one of the main advantages of using backoff protocols is that devices limit the number of times they attempt to access the shared resource (e.g., the channel), which is often the expensive operation. Previous jamming-resistant protocols [6, 44, 51–55], for example, are also energy efficient, using few broadcasts per packet, over sufficiently large time windows.

One observation that arises from looking at real radios is that *receiving messages is also expensive*. However, modern radios have the capacity to monitor the channel in a special low-power mode. This allows the device to determine if the channel is occupied or busy, but cannot distinguish between a message, a collision, or noise/jamming. Given such a mechanism, we can improve significantly on energy usage over prior protocols (e.g., [6, 44, 51–53], which requires that each device receive packets in order to distinguish between a single packet broadcast and a collision).

There are interesting connections between energy-efficiency and jamming. Since jamming requires that the radio be in the energy-expensive send state, there is a cost to the adversary for attacking.[3] Ideally, it should be more expensive for the adversary to attack than for the good devices to resist; this is the idea behind ***resource competitiveness*** (see [23, 24, 36]). More generally, jamming-resistant results include heuristics [1, 4, 5, 10, 38, 39, 43, 47, 60] and approaches with provable guarantees [9, 18, 19, 21, 22, 40, 45].

Results. We give an energy-efficient, dynamic, jamming-resistant, protocol NOISEOFF for a single-hop network, in which devices can send and receive packets over a shared wireless communication channel without the use of any global broadcast schedule, shared secrets, or centralized point of control. Guarantees are provided for both the case where the number of packets is finite and infinite.

We interpret jamming resistance to mean that in the event of a jamming attack, NOISEOFF (1) achieves good throughput, and (2) guarantees that a correct device uses far less energy than the adversary.

Our jamming resistance is consistent with the notion of ***resource competitiveness*** (see [23]) where, in our setting the resource is energy. Specifically, we show that: (i) NOISEOFF is energy efficient in the absence of jamming, and (ii) the additional cost incurred to each packet by jamming is exponentially smaller than the cost to the attacker for jamming (see [23] for more details). Thus, malfunctioning or malicious devices deplete their respective energy supplies much more rapidly than the correct devices.

Theorem 1. *Let J denote the total number of slots for which the adversary jams. For a finite number of packets n injected into the system, where n is chosen but not revealed* a priori *by the adversary,* NOISEOFF *has expected constant throughput*[4] *and an expected energy consumption of $O(\log^2(n+J))$ per packet.*

Extending these results to the infinite case, we show that both throughput and energy bounds hold true at a countably infinite number of slots over (countably infinite) execution time. This result is expressed by the following theorem. (It would be unreasonable to expect the bounds to apply at every step of the execution—that is not even typically true of finite backoff processes.)

Theorem 2. *For any time t, denote by J_t the number of slots jammed before t, and denote by η_t the maximum number of packets concurrently in the system before t. And suppose an infinite number of packets will be injected into the system. Then for any time s, there exists a time $t \geq s$ such that* NOISEOFF *has constant throughput with probability 1, and expected average energy consumption of $O(\log^2 \eta_t + \log^2 J_t)$ per packet.*

[3] In practice, jamming may be cheaper than sending and receiving since only a portion of a data slot need be corrupted for the adversary to be successful. However, activating the circuitry and onboard radio to conduct jamming is still expensive. Additionally, the use of smaller packet sizes (see [5]) further reduces any discrepancy. Our asymptotic results hold as long as the relative costs between jamming, sending, and receiving are proportional to within a (possibly large) constant.

[4] Constant throughput means that either a constant fraction of the slots are successful transmissions, or a constant fraction of the slots are jammed—see Section 2.

Approach. Here is a simple protocol, which NOISEOFF builds upon. When a packet arrives in the system, it monitors a special ***control channel*** to listen for a ***busy signal***.[5] Only when there is no busy signal does the packet become active and begin making broadcast attempts. Each active packet broadcasts continually on the control channel, which ensures that no newcomers become active. It is not hard to prove constant throughput if the active packets run an efficient batch backoff protocol (e.g., SawTooth Backoff [7]).

This basic protocol has abominable energy consumption, since every active packet continually broadcasts on the control channel. A natural optimization is to make active packets broadcast probabilistically on the control channel (with probability much less than 1). With this change, the occasional ***control failure*** may occur, where some packets are active but no message is sent on the control channel.

It remains to specify how the protocol copes with control failures. At a high level, there are two options: either the active packets continue, or they ***reset*** and join afresh with the next batch. Resetting has the advantage of synchronizing active packets (and simplifying the analysis). Moreover, if packets never reset, then the throughput can be lousy. But if packets reset too eagerly, then the energy bounds suffer.[6] Continuing can lead to better energy bounds, but allows packets to get out of sync. Jamming exacerbates the situation.

Our solution lies somewhere in the middle. Active packets do sometimes reset. But a packet has to wait until a large-enough constant fraction of the (data) slots have been empty during its lifetime before resetting. In principle, this means that the packet had plenty of chances to broadcast successfully in its lifetime before it ended up resetting instead.

There are several tricky aspects to the analysis. Most notably, there are likely to be concurrently active packets at different stages of their protocols. The contention created by the new packets may delay the old packets, and prevent good throughput. There is no fixed bound on how long it takes for a bolus of high contention to dissipate. We need to show that, nonetheless, the contention stays low enough, most of the time, to make progress. This means dealing with the adaptivity of the adversary, both in adding jobs and jamming, and that in infinite executions, we expect bad things to happen infinitely often.

2 Model

Multiple Access Channel. Time is discretized into ***slots***, where each slot is the length of one packet. When there is no transmission on the channel during a slot, we call that slot ***empty***. A slot is ***full*** when one or more packets are broadcast. When exactly one packet is broadcast in a slot, that packet transmits successful, and we call the slot ***successful***. When two or more packets are broadcast during the same slot, a ***collision*** occurs. When there is a collision, there is noise on the channel, but all packets transmitting are unsuccessful. We assume that a device transmitting a packet can determine whether its transmission is successful; this is a standard assumption in the backoff literature (for examples, see [27, 34, 37]), unlike the wireless setting where a full medium access control (MAC) protocol would address acknowledgements and many other issues. Here, our work focuses solely on the sending side (backoff component) of the problem.

In general, we assume that sending or receiving a packet costs one unit of energy. We also assume that devices have a low-power "monitoring" mode for monitoring the channel that can distinguish whether a slot is full or empty at zero cost. While performing low-power monitoring, a device does *not* actually receive any packets that successfully sent, nor can it distinguish between a successful slot and a collision.[7]

Finally, we do not assume a global clock, i.e., that packets agree on a universal numbering scheme for the slots. We do, however, assume that devices know the parity of each slot (whether it is an even or odd slot) as this simplifies the presentation. In Section 6, we discuss approaches to implement this assumption.

[5] Busy signals have also been employed in MAC protocols (see [32, 57] for examples) for coping with hidden terminal effects.

[6] There are instances where a packet's energy to be polynomial in the adversary's energy rather than polylogarithmic.

[7] Assessment of channel activity can be done via an "energy detector" (see [50]). Low-power monitoring can be cheaper than, for example, actually receiving a message. For example, cellphones employ a pilot channel to measure channel quality (i.e. often appearing as bars on the display) rather than engaging in the more expensive operation of a phone call. Low-power monitoring is technically feasible for other small wireless devices, and it is plausible that it will be increasingly common (see [56] for progress in this direction, where the cost of energy detection is reduced by a factor of more than 17).

Arbitrary Dynamic Packet Arrivals. Over time, new packets ***arrive***. These arrivals are arbitrary—we do not assume any bounded arrival rate. A packet is ***live*** at any time between its arrival and its successful transmission. The number of packets in the system may vary arbitrarily over time, and this number is unknown to all nodes. Without loss of generality, we assume that there is at least one packet alive during the entire execution of the protocol. (If not, we simply ignore any slot during which there are no live packets.)

Adversary. We postulate an ***adversary*** who governs two aspects of the system's dynamics. First, the adversary determines the (finite) number of new packets that arrive in the system in each slot. Second, the adversary may arbitrarily block, or ***jam***, slots by creating noise on the channel. This jamming causes a collision if any packet is being transmitted.

The adversary is adaptive with one exception—the adversary must decide *a priori* whether the execution will contain infinitely many packets or a finite number n of packets. In the finite case, the adversary chooses the number n *a priori*. The nodes themselves do not know whether the instance is infinite or finite, and in the finite case, do not know n. In all other ways, the adversary is adaptive: it may make all arrival and jamming decisions with full knowledge of the current and past system state; at the end of a given slot, the adversary learns everything that has happened in that slot. The adversary does not know the random choices that packets make in slot t prior to the end of slot t. If a slot is jammed, then all packets experience this jamming (all packets witness the same jamming schedule); that is, our adversary is 1-uniform (see [51]).

Throughput. We define the throughput for intervals, finite instances, and infinite instances; see [6, 51–53] for similar definitions. The throughput $\lambda \in [0, 1]$ of an interval I is the fraction of non-wasted slots in I:

Definition 3. *The **throughput of an interval** $\mathcal{I}$ with $N_{\mathcal{I}}$ successful transmissions and $J_{\mathcal{I}}$ jammed slots is $\lambda = (N_{\mathcal{I}} + J_{\mathcal{I}})/|\mathcal{I}|$.*

In the absence of jamming, throughput is simply the fraction of slots that contain successful transmissions. By contrast, a jammed slot is never seen as wasted, since it could never be used for a successful packet transmission. Notice that if only a $\lambda/2$ fraction of slots are jammed, then λ-throughput guarantees that a constant fraction of the non-jammed slots are used for successful transmission (i.e., as in [6, 51–53]). We next extend this definition to talk about the performance on finite and infinite instances.

Definition 4. *An instance with n packets, J jammed slots, and in which every packet completes by slot T, has throughput $\lambda = (n + J)/T$.*

Definition 5. *An infinite instance has λ-throughput if, for any slot t, there exists a slot $t' \geq t$ where interval $[0, t']$ has λ throughput.*

3 Algorithm

This section presents the details of our backoff protocol. We analyze its throughput in Section 4, and its energy usage in Section 5.

To simplify presentation, we assume throughout that there are two communication channels, a ***data channel***, on which packets are broadcast, and a ***control channel***, on which a "busy signal" is broadcast. It is straightforward to simulate these two channels with only one, using odd slots to simulate the data channel and even slots to simulate the control channel. (Recall that devices agree on slot parity.)

In each slot, packets may arrive in the system. For a given packet u, let s_u be the number of slots it has been active for. Our protocol has the following structure (see Figure 1 for pseudocode):

- Initially, each packet is ***inactive***; it makes no attempt to broadcast on either channel.
- Inactive packets monitor the control channel (in low-power mode). As soon as there is an empty slot on the control channel, the packet becomes ***active***.

NOISEOFF for a node u that has been active for s_u rounds

- With probability $\frac{c \max\{\ln s_u,\, 1\}}{s_u}$, send on the control channel
- With probability $\frac{d}{s_u}$ send m_u on the data channel and, if successful, then terminate.
- Monitor the data channel with low-power listening. If at least a $\gamma = 7/8$-fraction of data slots have been clear since node u's injection into the system, then become inactive.

NOISEOFF for an inactive node

- Monitor each control round with low-power listening. If a round is clear, then become active next round

Figure 1: Pseudocode for NOISEOFF.

- Once a packet is active, in every time slot it broadcasts on the data channel with some probability proportional to how long it has been active, i.e., a packet u broadcasts with probability d/s_u, for some constant $d \leq 1/2$. It also broadcasts on the control channel with probability $c \max(\ln s_u, 1)/s_u$, for some constant $c > 0$.
- The packet remains active until either it transmits successfully, or it sees too many empty slots. More specifically, if packet u has observed $(7/8)s_u$ empty slots (via low-power monitoring), the packet reverts to an ***inactive*** state, and the process repeats.

4 Throughput Analysis

In this section, we analyze the throughput of the NOISEOFF protocol, showing that it achieves constant throughput in both the finite and infinite cases. All omitted proofs appear in Appendix A.

Let s_j^t be the age of packet j in slot t, i.e., the number of slots that it has been active.

At time t, we define the ***contention*** to be $X(t) = \sum_i 1/s_j^t$, where we sum over all the active packets. Thus, the expected number of broadcasts on the data channel in slot t is $dX(t)$.

We now divide the packets that are active in slot t into ***young*** and ***old*** packets. For every slot t, we define time σ_t as the minimum time $\geq t$ where the following hold:

- $\sum_{i:0<s_i^t\leq\sigma_t} 1/s_i^t \geq X(t)/2$;
- $\sum_{i:s_i^t\geq\sigma_t} 1/s_i^t \geq X(t)/2$.

That is, the active packets with age $\leq \sigma_t$ have at least half the contention, and the active nodes with age $\geq \sigma_t$ have at least half the contention. We call the former set the ***young*** nodes and the latter set the ***old*** nodes. Note that some nodes (i.e., those with age exactly σ_t) may be *both* young and old.

Lemma 6. *For all times t, σ_t is well-defined.*

We say that a ***control failure*** occurs in slot t if no node broadcasts on the control channel during the slot. Recall that (1) a packet activation can occur only immediately after a control failure and (2) a packet j ***resets*** at time t if t is the first slot during j's lifetime $[t - s_j^t, t]$ of s_j slots, for which at least $\frac{7}{8}s_j$ slots were empty.

4.1 Individual Slot Calculations

First we look at the probability of a successful broadcast as a function of the contention. The next two lemmas explicitly require $d \leq 1/2$.

Lemma 7. *For a given slot t in which there is no adversarial jamming, the probability that some packet successfully broadcasts at time t is at least $\frac{d\,X(t)}{e^{2d\,X(t)}}$.*

Next, we give bounds on the probability of any broadcast, as a function of the contention.

Lemma 8. *The probability that some packet is broadcast (not necessarily successfully) in slot t is at least $1 - e^{-d\,X(t)}$ and at most $1 - e^{-2d\,X(t)}$. The probability of a collision in the slot is at most $(1 - e^{-2d\,X(t)})^2$.*

4.2 Epochs, Streaks, and Interstitial Slots

An execution is divided into two types of periods: ***epochs*** and ***interstitial slots***.

Definition 9. *Each time a packet is activated, a new **epoch** begins. An epoch is divided into **streaks**. The first streak begins with the beginning of the epoch. If a streak begins at time t, then it either ends at time $t + \sigma_t$ or when a new epoch begins—whichever happens first. If, when a streak ends at time t, the contention $X(t) \geq 16$, then a new streak begins. Otherwise, the streak and epoch ends. After an epoch ends, the interstitial slots begin. The interstitial slots continue until the next epoch begins (with another packet activation).*

In general, we say that an epoch is ***jammed*** if at least 1/4 of its slots are jammed. We next bound the change in contention during a streak. (By definition, there are no control failures and packet activations.)

Lemma 10. *Assume that some streak begins at time t and that no control failures occur during the streak. Then $X(t + \sigma_t) \leq 3X(t)/4$.*

Lemma 11. *Assume that some streak begins at time t, where $X(t) \geq 8$, and that no resets occur during the streak. Then for all $t' \in [t, t + \sigma_t]$, $X(t') \geq X(t)/8$.*

4.3 Control Failures

Next we look at the probability of a control failure, as a function of the contention. The lemma follows by showing that in each slot, the old packets provide sufficient contention throughout the streak to prevent a control failures.

Lemma 12. *For a fixed time t when a streak begins, consider a control slot at time t' during the streak beginning at time t. Assume that there are no control failures or resets during the streak prior to time t'. Then the probability of a control failure in slot t' is at most $(\sigma_t)^{-\frac{cX(t)}{2}}$.*

We can then take a union bound over the slots in the streak and conclude:

Lemma 13. *For a fixed time t, consider a streak beginning at time t. Assume that no reset occurs during the streak. The probability that a control failure occurs in the interval $[t, t + \sigma_t]$ is at most $(\sigma_t)^{-bX(t)}$ for a constant $b > 0$ depending only on constant c in our algorithm.*

4.4 Bounding Resets

We next bound the probability that a reset takes place during an epoch. If a packet has not reset by the time an epoch begins, we show that with constant probability, it does not reset during the epoch. This is true since for any prefix of the epoch, there are sufficiently many broadcasts to prevent a reset. We first look at an abstract coin flipping game:

Lemma 14. *Consider a sequence of Bernoulli trials each with probability p of success. If the first $16/p$ trials are all successful with probability at least q, then with probability at least $q/2$, for all i, the first i trials contain at least $ip/4$ successes.*

We now conclude in the next lemma that a reset occurs during an epoch with a bounded constant probability.

Lemma 15. *Consider an epoch that begins at time t. Then a reset occurs during the epoch with probability at most $1 - (q/2)$ for some constant $q > 0$ depending only on d.*

4.5 Successful Streaks

The notation $S(t)$ refers to a streak beginning at time t and continuing for σ_t slots. A streak is ***successful*** if there are no resets or packet activations during the streak. Notice that during a successful streak, we know that the contention is always at least $X(t)/8$ (by Lemma 11), and that at the end of the streak it is at most $3X(t)/4$ (by Lemma 10). We say that successful streaks $S(t)$ and $S(t')$ are consecutive if $t' = t + \sigma_t$.

Lemma 16. *Let $S(t)$ and $S(t')$ be consecutive successful streaks and assume contention is non-zero over both streaks. Then, $\sigma_{t'} > \sigma_t$.*

We can then show that, with constant probability, an epoch is "successful." This is in many ways the key result that allows to later show constant throughput, and it follows from an analysis of the change in contention and a union bound over the streaks.

Lemma 17. *For an epoch beginning at time t, with constant probability, every streak is successful.*

We say that an epoch is *jammed* if at least $1/4$ the slots in the epoch are jammed. The following corollary shows that we get constant throughput in an epoch with constant probability.

Corollary 18. *For an epoch beginning at time t and of length T, with constant probability: (i) every streak in the epoch is successful; (ii) the last streak is of length at least $T/2$; (iii) the contention throughout the last streak is between 1 and 256; and (iv) if the epoch is unjammed, then at least $\Omega(T)$ packets are broadcast.*

Proof (sketch). Conclusion (i) follows from Lemma 17, while conclusion (ii) follows from Lemma 16. Conclusion (iii) follows because there are no resets or packet activations, and hence the contention decreases by at most a factor of 16; however, since the epoch ends, we know that it is no greater than 16 when the last streak ends. Conclusion (iv) follows from observing that, in an unjammed slot, there is a constant probability that a packet is broadcast, and a constant probability that one is not (due to an empty slot or a collision). If at most $1/4$ of the epoch is jammed, then at most half of the slots in the last streak are jammed, and of these $T/4$ unjammed slots in the last epoch, in expectation, only a constant fraction are *not* successful broadcasts. Thus, by Markov's inequality, with constant probability, at most a constant fraction of these $T/4$ slots are not successful broadcasts, and hence with constant probability, at least $\Omega(T)$ packets are broadcast. □

4.6 Bad Borrower Game

We have shown that each epoch is ***good*** (satisfying Corollary 18) with constant probability. We now abstract away some details, defining a simple game between two players: the ***lender*** and the ***borrower***. There are two key parameters: a probability p and a fraction $\alpha \in (0, 1)$. The game proceeds in iterations, where in each, the borrower borrows an arbitrary (adversarially chosen) amount of money from the lender, at least one dollar. With probability p, at the end of the iteration, the borrower repays a fraction α of the money.

The correspondence to our situation is as follows: each iteration is associated with an unjammed epoch, the length of the epoch defines the money borrowed, and the number of successful broadcasts defines the money repaid. In a good epoch, which occurs with constant probability p, if the epoch is unjammed, then we get constant throughput and hence the borrower is repaid an α fraction of her money. In a bad epoch, by contrast, we allow for the worst case, which is no money paid back at all (no throughput).

For the ***finite Bad Borrower game***, there is a predetermined maximum amount that the lender can be repaid: after the borrower has been repaid n dollars, the game ends. This corresponds to a finite adversary that injects exactly n packets. In the ***infinite Bad Borrower game***, the game continues forever, and an infinite amount of money is lent. This corresponds to infinite instances, where the adversary injects packets forever.

4.7 Finite Bad Borrower Game

Our goal in this section is to show that, when the borrower has repaid n dollars, he has borrowed at most $O(n)$ dollars. This corresponds to showing that n packets are successfully broadcast in $O(n)$ time, ignoring the interstitial slots (which we will come back to later).

We assume throughout this section that n is the maximum amount of money repaid throughout the game, i.e., the adversary injects n packets in an execution. There is a simple correspondence lemma which bounds the amount of money that the borrower can borrow:

Lemma 19. *In every iteration of the finite bad borrower game, the borrower borrows at least one dollar and at most n dollars.*

We now argue, via an analysis of the expected repayments, that when the finite bad borrower game ends, the expected cost to the lenders is $O(n)$:

Lemma 20. *Over the entire execution of the bad borrower game, the expected number of dollars borrowed, in total, is $O(n)$.*

4.8 Infinite Bad Borrower Game

In order to analyze an infinite executions, with injections that continue forever, we look at the infinite bad borrower game. Recall that at any given point in the game, for parameter k chosen in advance, if there have been k dollars borrowed up to that point, then in expectation there have been $O(kp\alpha)$ dollars repaid. In the infinite case, we can conclude something stronger: there are an infinite number of times at which the borrower has repaid at least an $\alpha/2$ fraction of the total dollars borrowed.

Lemma 21. *For all iterations r of the infinite bad borrower game, there is some iteration $r' > r$ such that if the lender has lent k dollars through slot r', then the borrower has repaid at least $kp\alpha/2$ dollars, with probability 1.*

As a corollary, if we only consider the epochs, ignoring the contribution from the interstitial slots, we can show constant throughput for infinite executions. Specifically, we can use the infinite bad borrower game to define "measurement points," thus showing that in an infinite execution, there are an infinite number of points at which we get constant throughput (if we ignore the contribution from the interstitial slots).

Corollary 22. *If we take an infinite execution and remove all slots that are not part of an epoch, then the resulting execution has constant throughput.*

We later show (Subsection 4.10) that the contribution from the interstitial slots does not hurt the throughput, meaning that we get constant throughput taking into account all slots.

4.9 Interstitial Slots and Expected Throughput for Finite Instances

We begin by considering the finite case where there are n packets injected in total. We bound the length of the interstitial slots, after which we are able to then prove expected constant throughput. A ***full reset*** at time slot s refers to the case where all packets that are active in slot s restart in slot $s+1$. The next lemma proves we may view the execution of the algorithm as broken into periods separated by full resets. We refer to each such period as a ***reset interval***.

Lemma 23. *If an active packet u that becomes active in slot t_0 and resets in slot $t > t_0$. Then, all active packets injected in $[t_0, t]$ reset by slot t with probability 1.*

We can now make a key observation: at all times at most $7/8$ of the slots are empty. This follows immediately from Lemma 23, from an analysis of the reset intervals.

Lemma 24. *With probability 1, the following hold:*

- *Any reset interval consists of at most a $\frac{7}{8}$-fraction of empty slots.*
- *Any interval of slots after a full reset consists of at most a $\frac{7}{8}$-fraction of empty slots.*
- *If there are no reset intervals, then the slots from when the first packet is broadcast and the last packet succeeds consists of at most a $\frac{7}{8}$-fraction of empty slots.*

We observe that the empty slots cannot hurt the throughput by more than a constant factor overall. This is because, by Lemma 24, any prefix of the execution contains at most a $7/8$ fraction of empty slots. Hence if the execution has at most $O(n + J)$ non-empty slots, it also has at most $O(n + J)$ empty slots.

We now analyze the throughput for the non-empty interstitial slots. First, we know that at most J of these slots are jammed. We now bound the number of non-empty, non-jammed interstitial slots. For each such slot, there is at least a constant probability of a successful transmission. Thus after $O(n)$ such slots, in expectation, all the packets have broadcast.

Lemma 25. *For a slot s, let $e_{\geq 2}$ denote the event where two or more packets are broadcast in s, and let $e_{=1}$ denote the event where one packet is broadcast in slot s. If s is an interstitial slot, then $Pr(e_{\geq 2}) = O(Pr(e_{=1}))$.*

Lemma 26. *There are at most $O(n)$ full, unjammed interstitial, slots in expectation.*

We can now prove our claim in Theorem 1 regarding expected constant throughput:

Lemma 27. *If the adversary injects a total of n packets,* NOISEOFF *achieves expected constant throughput.*

4.10 Interstitial Slots for Infinite Instances

We now show that the contribution from the interstitial slots does not hurt the throughput in infinite executions. To do so, we deterministically bound the contribution from the empty interstitial slots. We show that as long as we pick "measurement points" that are sufficiently large that from then on the non-empty interstitial slots do not hurt. We use the following well known facts about random walks [42].

Fact 28. *Suppose that we have a biased random walk on a line with fixed step size, where the probability of going right is at least p, the probability of going left is at most $1-p$, and the step size right is δ_r and the step size left is δ_ℓ. Suppose that $p\delta_r > (1-p)\delta_\ell$. Then if the random walk starts at the origin, the probability of returning to the original is some constant strictly less than 1.*

Corollary 29. *For any such biased random walk, there there is a last time that the walk returns to the origin.*

We use Corollary 29 to bound the ratio of collisions to broadcasts in unjammed, non-empty interstitial rounds. Notably, from some point on, the number of collisions is always at most a constant factor of the number of broadcasts, and hence yields constant throughput. We also observe that the empty slots cannot hurt the throughput by more than a constant factor overall, because, by Lemma 24, any prefix of the execution contains at most a $7/8$ fraction of empty slots. Finally, we can bound the jammed interstitial slots by J. From this, we conclude that we achieve constant throughput, as claimed in Theorem 2:

Lemma 30. *In an infinite execution,* NOISEOFF *achieves constant throughput.*

5 Energy Analysis

In this section, we analyze the energy usage of the backoff protocol. Since the protocol never receives any messages, and since low-power monitoring is free, we have only to analyze the number of times each packet broadcasts. Our goal is to show that in the absence of jamming, the number of broadcasts is small, and that the adversary requires a significant amount of jamming to cause even a small increase in the energy usage.

We first analyze how often a packet resets, showing that it is likely to succeed before it has a chance to reset. This ensures that a packet cannot be forced to spend a large amount of energy via repeated resets.

Lemma 31. *With constant probability, a packet succeeds before it resets.*

From this we can immediately bound the distribution on the number of times a packet resets.

Corollary 32. *For any positive integer k, the probability that a packet resets k times is at most $1/e^{\Theta(k)}$.*

We can now bound the total energy that a packet spends during the first t slots after its arrival. If it has not yet reset by time t, it is easy to see that it has spent $O(\log^2 t)$ energy, in expectation—and we have show above that a packet is unlikely to reset too many times. This yields:

Lemma 33. *In the first t slots following a packet's arrival, it expends $O(\log^2(t))$ energy in expectation.*

We can now prove our claim in Theorem 1 regarding the expected energy consumption per packet by observing that we have already bounded the expected length of the execution in Lemma 27:

Lemma 34. *Consider the finite case, let n be the number of injected packets, and let J be the number of jammed slots. Then the expected energy each packet expends is $O(\log^2(n+J))$.*

We can now prove our claim in Theorem 2 regarding the expected energy consumption per packet. The proof examines separately the jammed slots, the unjammed slots with young packets, and the unjammed packets with old packets—in each case, we show that the energy used is bounded.

Lemma 35. *Consider any time t in the infinite case at which we have λ throughput, for constant λ. Let J_t be the total number of jammed slots before t, and let η_t be the maximum contention prior to time t. Then the expected average energy per packet is $O(\log^2(\eta_t) + \log^2(J_t))$.*

6 Conclusion

We presented a simple backoff protocol that provides constant throughput with online packet arrivals. It tolerates disruptive noise/jamming, and by leveraging low-power monitoring, it ensures low energy usage.

One issue not discussed earlier is "cohabiting" (or "co-existing") networks, when different applications are sharing the same communication channel. If these applications are unaware of each other (e.g., treating messages from the other application as noise, or as sent by an adversary), then existing backoff protocols and MAC layers (e.g., [6]) no longer guarantee good throughput. By contrast, since the algorithm here has no sensitivity to which messages are received—only monitoring for when the channel is occupied or free—it will still guarantee constant throughput, collectively, in such settings. (See, for example [53], which also addresses the problem of co-existing networks.)

Also, we assumed throughout the paper that packets agree on the parity of the slot, allowing us to simulate the control and data channels on a single channel. We believe that this parity assumption can be dropped using the following trick to synchronize packets. Most of the time the control slots and data slots alternate. The broadcast probabilities are correlated so that whenever a packet transmits in the data slot, it transmits in the corresponding (preceding) control slot. However, whenever any control slot is empty, the next slot becomes an additional control slot. This repetition means that despite any jamming, any new packet arriving in the system can identify an empty control slot so that it can become active. Only a small modification to the throughput analysis is required.

Another technical issue is whether we can tolerate less predictable noise, i.e., a more powerful adversary that can cause different packets to observe different conditions (i.e., an n-uniform adversary). Such a model better captures a multi-hop network where noise may come from different directions. The key challenge is that packets may see different slots as empty, and hence reset at different times. While resets are further desynchronized, it remains the case that there cannot be too many truly empty slots (i.e., empty for everyone) without causing resets. Hence again, we suspect that the results of this paper hold, with small changes.

The larger question is to examine how NOISEOFF operates in a multi-hop network. Previous works on jamming resistance [55] have shown that it is possible to design a jamming-resistant backoff protocol that continues to guarantee good throughput, even when the devices sending the packets are distributed over a large area. Other open questions include looking at a reactive adversary (that can decide on a jamming strategy after learning what the packets will do in a given slot), and looking at specific models of radio communication (e.g., the SINR model).

References

[1] Ghada Alnifie and Robert Simon. A Multi-Channel Defense Against Jamming Attacks in Wireless Sensor Networks. In *Proceedings of the 3rd ACM Workshop on QoS and Security for Wireless and Mobile Networks*, pages 95–104, 2007.

[2] Lakshmi Anantharamu, Bogdan S. Chlebus, and Mariusz A. Rokicki. Adversarial Multiple Access Channel with Individual Injection Rates. In *Proceedings of the 13th International Conference on Principles of Distributed Systems (OPODIS)*, pages 174–188, 2009.

[3] Damal Kandadai Arvind, Khaled Elgaid, Thomas Krauss, Allan Paterson, Robert Stewart, and Iain Thayne. Towards an Integrated Design Approach to Specknets. In *Proceedings of the IEEE International Conference on Communications (ICC)*, pages 3319–3324, 2007.

[4] Nils Aschenbruck, Elmar Gerhards-Padilla, and Peter Martini. Simulative Evaluation of Adaptive Jamming Detection in Wireless Multi-hop Networks. In *International Conference on Distributed Computing Systems Workshops*, pages 213–220, 2010.

[5] Farhana Ashraf, Yih-Chun Hu, and Robin Kravets. Demo: Bankrupting the Jammer. In *Proceedings of the* 9^{th} *International Conference on Mobile Systems, Applications, and Services (MobiSys)*, 2011.

[6] Baruch Awerbuch, Andrea Richa, and Christian Scheideler. A Jamming-Resistant MAC Protocol for Single-Hop Wireless Networks. In *Proceedings of the* 27^{th} *ACM Symposium on Principles of Distributed Computing (PODC)*, pages 45–54, 2008.

[7] Michael A. Bender, Martin Farach-Colton, Simai He, Bradley C. Kuszmaul, and Charles E. Leiserson. Adversarial contention resolution for simple channels. In *Proc. 17th Annual ACM Symposium on Parallelism in Algorithms and Architectures (SPAA)*, pages 325–332, 2005.

[8] Michael A. Bender, Jeremy T. Fineman, and Seth Gilbert. Contention resolution with heterogeneous job sizes. In *Proc. 14th Annual European Symposium on Algorithms (ESA)*, pages 112–123, 2006.

[9] Vartika Bhandhari, Jonathan Katz, Chiu-Yuen Koo, and Nitin Vaidya. Reliable Broadcast in Radio Networks: The Bounded Collision Case. In *Proceedings of the ACM Symposium on Principles of Distributed Computing (PODC)*, pages 258 – 264, 2006.

[10] Timothy Brown, Jesse James, and Amita Sethi. Jamming and Sensing of Encrypted Wireless Ad Hoc Networks. In *Proceedings of the* 7^{th} *ACM International Symposium on Mobile Ad Hoc Networking and Computing*, pages 120–130, 2006.

[11] Bogdan S. Chlebus, Leszek Gasieniec, Dariusz R. Kowalski, and Tomasz Radzik. On the Wake-up Problem in Radio Networks. In *Proceedings of the 32nd International Colloquium on Automata, Languages and Programming (ICALP)*, pages 347–359, 2005.

[12] Bogdan S. Chlebus and Dariusz R. Kowalski. A Better Wake-up in Radio Networks. In *Proceedings of 23rd ACM Symposium on Principles of Distributed Computing (PODC)*, pages 266–274, 2004.

[13] Bogdan S. Chlebus, Dariusz R. Kowalski, and Mariusz A. Rokicki. Adversarial queuing on the multiple-access channel. In *Proc. Twenty-Fifth Annual ACM Symposium on Principles of Distributed Computing (PODC*, pages 92–101, 2006.

[14] Bogdan S. Chlebus, Dariusz R. Kowalski, and Mariusz A. Rokicki. Adversarial queuing on the multiple access channel. *ACM Transactions on Algorithms*, 8(1):5, 2012.

[15] Marek Chrobak, Leszek Gasieniec, and Dariusz R. Kowalski. The Wake-up Problem in Multihop Radio Networks. *SIAM Journal on Computing*, 36(5):1453–1471, 2007.

[16] Crossbow. MICAz Wireless Measurement System. `http://www.openautomation.net/uploadsproducto`

[17] Crossbow. MICA2 Wireless Measurement System. `http://www.eol.ucar.edu/isf/facilities/isa/ir`

[18] Shlomi Dolev, Seth Gilbert, Rachid Guerraoui, and Calvin Newport. Gossiping in a Multi-channel Radio Network: An Oblivious Approach to Coping with Malicious Interference. In *Proceedings of the International Symposium on Distributed Computing (DISC)*, pages 208–222, 2007.

[19] Shlomi Dolev, Seth Gilbert, Rachid Guerraoui, and Calvin Newport. Secure Communication over Radio Channels. In *Proceedings of the Symposium on Principles of Distributed Computing (PODC)*, pages 105–114, 2008.

[20] Mihály Geréb-Graus and Thanasis Tsantilas. Efficient optical communication in parallel computers. In *Proc. of the 4th Annual Symposium on Parallel Algorithms and Architectures*, pages 41–48, 1992.

[21] Seth Gilbert, Rachid Guerraoui, Dariusz Kowalski, and Calvin Newport. Interference-Resilient Information Exchange. In *INFOCOM*, pages 2249–2257, 2009.

[22] Seth Gilbert, Rachid Guerraoui, and Calvin C. Newport. Of Malicious Motes and Suspicious Sensors: On the Efficiency of Malicious Interference in Wireless Networks. In *International Conference On Principles Of Distributed Systems (OPODIS)*, pages 215–229, 2006.

[23] Seth Gilbert, Valerie King, Jared Saia, and Maxwell Young. Resource-Competitive Analysis: A New Perspective on Attack-Resistant Distributed Computing. In *Proceedings of the* 8^{th} *ACM International Workshop on Foundations of Mobile Computing*, 2012.

[24] Seth Gilbert and Maxwell Young. Making Evildoers Pay: Resource-Competitive Broadcast in Sensor Networks. In *Proceedings of the* 31^{th} *Symposium on Principles of Distributed Computing (PODC)*, pages 145–154, 2012.

[25] Leslie Ann Goldberg, Mark Jerrum, Tom Leighton, and Satish Rao. Doubly logarithmic communication algorithms for optical-communication parallel computers. *SIAM Journal on Computing*, 26(4):1100–1119, August 1997.

[26] Leslie Ann Goldberg and Philip D. MacKenzie. Analysis of practical backoff protocols for contention resolution with multiple servers. In *SODA'96*, pages 554–563, 1996.

[27] Leslie Ann Goldberg, Philip D. Mackenzie, Mike Paterson, and Aravind Srinivasan. Contention Resolution with Constant Expected Delay. *Journal of the ACM*, 47(6):1048–1096, 2000.

[28] Leslie Ann Goldberg, Yossi Matias, and Satish Rao. An optical simulation of shared memory. *SIAM Journal on Computing*, 28(5):1829–1847, October 1999.

[29] Jonathan Goodman, Albert G. Greenberg, Neal Madras, and Peter March. Stability of binary exponential backoff. *Journal of the ACM*, 35(3):579–602, July 1988.

[30] Albert G. Greenberg, Philippe Flajolet, and Richard E. Ladner. Estimating the multiplicities of conflicts to speed their resolution in multiple access channels. *JACM*, 34(2):289–325, April 1987.

[31] Albert G. Greenberg and Shmuel Winograd. A lower bound on the time needed in the worst case to resolve conflicts deterministically in multiple access channels. *JACM*, 32(3):589–596, July 1985.

[32] Zygmunt J. Haas and Jing Deng. Dual Busy Tone Multiple Access (DBTMA) - A Multiple Access Control Scheme for Ad Hoc Networks. *Communications, IEEE Transactions on*, 50(6):975–985, Jun 2002.

[33] Johan Hastad, Tom Leighton, and Brian Rogoff. Analysis of backoff protocols for multiple access channels. In *STOC'87*, pages 241–253, New York, New York, May 1987.

[34] Johan Hastad, Tom Leighton, and Brian Rogoff. Analysis of Backoff Protocols for Mulitiple Access Channels. *SIAM Journal on Computing*, 25(4):1996, 740-774.

[35] Maurice Herlihy and J. Eliot B. Moss. Transactional memory: Architectural support for lock-free data structures. In *Proc. of the 20th Intnl. Conference on Computer Architecture.*, pages 289–300, San Diego, California, 1993.

[36] Valerie King, Jared Saia, and Maxwell Young. Conflict on a Communication Channel. In *Proceedings of the* 30^{th} *Symposium on Principles of Distributed Computing (PODC)*, pages 277–286, 2011.

[37] Byung-Jae Kwak, Nah-Oak Song, and Leonard E. Miller. Performance Analysis of Exponential Back-off. *IEEE/ACM Transactions on Networking*, 13(2):2005, 343-355.

[38] Mingyan Li, Iordanis Koutsopoulos, and Radha Poovendran. Optimal Jamming Attacks and Network Defense Policies in Wireless Sensor Networks. In *INFOCOM*, pages 1307–1315, 2007.

[39] Xin Liu, Guevara Noubir, Ravi Sundaram, and San Tan. SPREAD: Foiling Smart Jammers using Multi-layer Agility. In *Proceedings of the International Conference on Computer Communications (INFOCOM)*, pages 2536–2540, 2007.

[40] Dominic Meier, Yvonne Anne Pignolet, Stefan Schmid, and Roger Wattenhofer. Speed Dating Despite Jammers. In *Proceedings of the International Conference on Distributed Computing in Sensor Systems (DCOSS)*, pages 1–14, 2009.

[41] Robert M. Metcalfe and David R. Boggs. Ethernet: Distributed packet switching for local computer networks. *CACM*, 19(7):395–404, July 1976.

[42] Michael Mitzenmacher and Eli Upfal. *Probability and Computing: Randomized Algorithms and Probabilistic Analysis*. Cambridge University Press, 2005.

[43] Vishnu Navda, Aniruddha Bohra, Samrat Ganguly, and Dan Rubenstein. Using Channel Hopping to Increase 802.11 Resilience to Jamming Attacks. In *INFOCOM*, pages 2526–2530, 2007.

[44] Adrian Ogierman, Andrea Richa, Christian Scheideler, Stefan Schmid, and Jin Zhang. Competitive MAC under Adversarial SINR, 2013. `http://arxiv.org/abs/1307.7231`.

[45] Andrzej Pelc and David Peleg. Feasibility and Complexity of Broadcasting with Random Transmission Failures. In *Proceedings of the ACM Symposium on Principles of Distributed Computing (PODC)*, pages 334–341, 2005.

[46] Joseph Polastre, Robert Szewczyk, and David Culler. Telos: Enabling Ultra-Low Power Wireless Research. In *IPSN*, 2005.

[47] Christina Pöpper, Mario Strasser, and Srdjan Čapkun. Jamming-Resistant Broadcast Communication Without Shared Keys. In *Proceedings of the* 18^{th} *USENIX Security Symposium*, 2009.

[48] Prabhakar Raghavan and Eli Upfal. Stochastic contention resolution with short delays. *SIAM Journal on Computing*, 28(2):709–719, April 1999.

[49] Ravi Rajwar and James R. Goodman. Speculative lock elision: Enabling highly concurrent multithreaded execution. In *Proc. of the 34th Annual Intnl. Symposium on Microarchitecture*, pages 294–305, Austin, Texas, December 2001.

[50] Iyappan Ramachandran and Sumit Roy. On the impact of clear channel assessment on mac performance. In *IEEE Global Telecommunications Conference (GLOBECOM)*, pages 1–5, 2006.

[51] Andrea Richa, Christian Scheideler, Stefan Schmid, and Jin Zhang. A Jamming-Resistant MAC Protocol for Multi-Hop Wireless Networks. In *Proceedings of the International Symposium on Distributed Computing (DISC)*, pages 179–193, 2010.

[52] Andrea Richa, Christian Scheideler, Stefan Schmid, and Jin Zhang. Competitive and Fair Medium Access Despite Reactive Jamming. In *Proceedings of the* 31^{st} *International Conference on Distributed Computing Systems (ICDCS)*, pages 507–516, 2011.

[53] Andrea Richa, Christian Scheideler, Stefan Schmid, and Jin Zhang. Competitive and Fair Throughput for Co-Existing Networks Under Adversarial Interference. In *Proceedings of the* 31^{st} *ACM Symposium on Principles of Distributed Computing (PODC)*, 2012.

[54] Andrea Richa, Christian Scheideler, Stefan Schmid, and Jin Zhang. An Efficient and Fair MAC Protocol Robust to Reactive Interference. *IEEE/ACM Transactions on Networking*, 21(1):760–771, 2013.

[55] Andrea Richa, Christian Scheideler, Stefan Schmid, and Jin Zhang. Competitive Throughput in Multi-Hop Wireless Networks Despite Adaptive Jamming. *Distributed Computing*, 26(3):159–171, 2013.

[56] Mo Sha, Gregory Hackmann, and Chenyang Lu. Energy-efficient Low Power Listening for Wireless Sensor Networks in Noisy Environments. In *Proceedings of the* 12^{th} *International Conference on Information Processing in Sensor Networks (IPSN)*, pages 277–288, 2013.

[57] Cheng shong Wu and Victor O.K. Li. Receiver-Initiated Busy-tone Multiple Access in Packet Radio Networks. In *Proceedings of the ACM Workshop on Frontiers in Computer Communications Technology*, SIGCOMM '87, pages 336–342, 1988.

[58] John Paul Walters, Zhengqiang Liang, Weisong Shi, and Vipin Chaudhary. *Security in Distributed, Grid, Mobile, and Pervasive Computing. Chapter 17: Wireless Sensor Network Security: A Survey*. Auerbach Publications, 2007.

[59] Dan E. Willard. Log-logarithmic protocols for resolving ethernet and semaphore conflicts. In *STOC'84*, pages 512–521, Washington, D.C., April 30–May 2 1984.

[60] Wenyuan Xu, Wade Trappe, Yanyong Zhang, and Timothy Wood. The Feasibility of Launching and Detecting Jamming Attacks in Wireless Networks. In *MobiHoc*, pages 46–57, 2005.

A Omitted Proofs

Preliminaries

Lemma 6. *For all times t, σ_t is well-defined.*

Proof. Sort the packets by age so that $s_1 \leq s_2 \leq \ldots$. Let k be the minimum index such that $X(t)/2 \leq \sum_{i=1}^{k} 1/s_i \leq X(t)/2 + 1/s_k$. Let $\sigma_t = s_k$. Notice that the young packets have contention at least $X(t)/2$, and the old packets have contention at least $X(t)/2$. □

Individual Slot Bounds

Lemma 7. *For a given slot t in which there is no adversarial jamming, the probability that some packet successfully broadcasts at time t is at least $\frac{d\,X(t)}{e^{2dX(t)}}$.*

Proof. Packet j is successful with probability $\frac{d}{s_j}\prod_{i\neq j}(1-\frac{d}{s_i})$. At most one packet is successful, so the success events for each node are disjoint. The probability that some packet succeeds is thus at least $\frac{d}{s_1}\prod_i(1-\frac{d}{s_i}) + \frac{d}{s_2}\prod_i(1-\frac{d}{s_i}) + \cdots + \frac{d}{s_k}\prod_i(1-\frac{d}{s_i}) = \left(\prod_i(1-\frac{d}{s_i})\right)\cdot\sum_j \frac{d}{s_j} \geq \frac{dX(t)}{e^{2dX(t)}}$. The denominator follows from the fact that $1 - x \geq e^{-2x}$ for $0 \leq x \leq 1/2$, and hence $\prod_i(1-\frac{d}{s_i}) \geq e^{-2d\sum_i(1/s_i)}$. □

Lemma 8. *The probability that some packet is broadcast (not necessarily successfully) in slot t is at least $1 - e^{-d\,X(t)}$ and at most $1 - e^{-2d\,X(t)}$. The probability of a collision in the slot is at most $(1 - e^{-2d\,X(t)})^2$.*

Proof. The probability that no nodes broadcast is $\prod_i\left(1-\frac{d}{s_i}\right) \leq e^{-d\,X(t)}$. Conversely, $\prod_i\left(1-\frac{d}{s_i}\right) \geq e^{-2d\,X(t)}$ by the fact that $1 - x \geq e^{-2x}$ for $0 \leq x \leq 1/2$. The probability of a collisions is at most the square of the probability of a given broadcast, because this is the probability you get if we allow each packet to broadcast twice. □

Contention Bounds

Lemma 10. *Assume that some streak begins at time t and that no control failures occur during the streak. Then $X(t + \sigma_t) \leq 3X(t)/4$.*

Proof. During the streak, all the young packets at time t *at least* double in age (since they each have age at most σ_t), leading their contention to at least halve. Moreover, the young packets at time t have contention at least $X(t)/2$, so the total contention reduces by at least $X(t)/4$. Since there are no control failures, there are no new packets activated and hence no increase in contention. □

Lemma 11. *Assume that some streak begins at time t, where $X(t) \geq 8$, and that no resets occur during the streak. Then for all $t' \in [t, t + \sigma_t]$, $X(t') \geq X(t)/8$.*

Proof. Since there are no resets (and no activations, by definition) during the streak, the contention only decreases due to packets completing and due to increasing age. Consider the old packets at time t (which have contention at least $X(t)/2$ at time t). Since each of these packets at most doubles in age (since they have age $\geq \sigma_t$), their total contention remains at least $X(t)/4$ throughout the streak.

Some of these packets may finish, thus reducing the contention further. Assume that the old packet with the largest contention completes in every slot—notice that each such packet that finishes reduces the contention by at most $1/\sigma_t$. Thus, if one such packet finishes in each slot of the streak, the total contention is reduced by at most $\sigma_t/\sigma_t = 1$. Thus, throughout the streak, the total contention remains at least $X(t)/4 - 1 \geq X(t)/8$ (since $X(t) \geq 8$). □

Control Failures

Lemma 12. *For a fixed time t when a streak begins, consider a control slot at time t' during the streak beginning at time t. Assume that there are no control failures or resets during the streak prior to time t'. Then the probability of a control failure in slot t' is at most $(\sigma_t)^{-\frac{cX(t)}{2}}$.*

Proof. The probability of a control failure in slot t' is at most:

$$\begin{aligned}\prod_i \left(1 - \frac{c \ln s_i^{t'}}{s_i^{t'}}\right) &\leq e^{-c\sum_i \frac{1}{s_i^{t'}} \ln s_i^{t'}} \\ &= \sigma_t^{-c\sum_i \frac{\ln s_i^{t'}}{s_i^{t'} \ln \sigma_i}} \\ &\leq \sigma_t^{-c\sum_{\text{“old”}\, i} \frac{1}{s_i^{t'}}} \\ &\leq \sigma_t^{-cX(t)/2}\end{aligned}$$

where the probability on the last line is a function of t (not t') by Lemma 6. □

Lemma 13. *For a fixed time t, consider a streak beginning at time t. Assume that no reset occurs during the streak. The probability that a control failure occurs in the interval $[t, t + \sigma_t]$ is at most $(\sigma_t)^{-bX(t)}$ for a constant $b > 0$ depending only on constant c in our algorithm.*

Proof. Assuming there are no control failures during time $[t, t + s - 1]$, the probability of a control failure at time $t + s$ is at most $(\sigma_t)^{-\frac{cX(t)}{2}}$. Taking a union bound over the σ_t time slots, the probability of a control happening in any time slot is at most $(\sigma_t)^{-\frac{cX(t)}{2}+1}$. □

Bounding Resets

Lemma 14. *Consider a sequence of Bernoulli trials each with probability p of success. If the first $16/p$ trials are all successful with probability at least q, then with probability at least $q/2$, for all i, the first i trials contain at least $ip/4$ successes.*

Proof. Break up the trials into geometrically increasing subsequences of 2^k trials each. We say that a "failure" occurs in the kth subsequence if there are fewer than $p2^k/2$ successful trials within that subsequence. Using a Chernoff bound, the failure probability is at most $e^{-p2^k/8}$. Using a union bound, the probability of any failure for subsequence $k \geq \lg(1/p) + 4$ is at most $\sum_{k=\lg(1/p)+4}^{\infty} e^{-p2^k/8} = \sum_{j=1}^{\infty} e^{-2^j} \leq 1/2$. Therefore, with probability at least $q \cdot (1/2)$, the first $16/p$ trials are a success *and* every subsequence has at least $p2^k/2$ successes.

Now, suppose there is no failure in any subsequence, i.e., each has at least $p2^k/2$ successful trials. Pick any cutoff point in the subsequence of size 2^i and examine the total of 2^{i+1} trials trials up to this point. The previous subsequences for $k = 1, \ldots, i-1$ each contain at least $p2^k/2$ successful trials, for a total of at least $p \cdot 2^i/2$ successful trials. Therefore, up to the cutoff point, at least a $p/4$-fraction of the trials are successful. □

Lemma 15. *Consider an epoch that begins at time t. Then a reset occurs during the epoch with probability at most $1 - (q/2)$ for some constant $q > 0$ depending only on d.*

Proof. Imagine, for the sake of the proof, we flip coins with probability of heads (corresponding to a full slot) $p = 1 - e^{-d}$ for each slot of the epoch in advance. The sending probability for a packet u in a data slot is d/s_u for some constant $d \leq 1/2$. The probability of sending in each of the ***initial*** $16/p$ slots is

$d\prod_{j=1}^{16/p}(1/j) = \frac{d}{(16/p)!} = \Theta(1)$ and denote this (small) probability by q. By Lemma 14, with probability at least $q/2$, the number of heads in any prefix of size i is at least $ip/4$. Consider the case where this good event occurs (i.e., with probability $q/2$).

Now consider executing the protocol. Consider some slot t', assuming that there has been no reset in the epoch prior to t'. In slot t', as long as there have been no prior resets, we know that $X(t') \geq X(t)/8$, by Lemma 11. Thus, by Lemma 8, a slot is empty with probability at most e^{-d}, i.e., it is full with probability at least p. Thus, time slot t' is empty only if the coin flip for slot t' is a tails. This implies that since there are at least $(t'-t)p/4$ heads in the interval from t to t', there is no reset in slot t'. Continuing inductively, we conclude that if the initial coin flips are good, then there is no reset during the epoch. □

Streaks

Lemma 16. *Let $S(t)$ and $S(t')$ be consecutive successful streaks and assume contention is non-zero over both streaks. Then, $\sigma_{t'} > \sigma_t$.*

Proof. Starting from slot t, let A denote the set of active packets after σ_t slots. Then, since the contention is non-zero, there exist remaining active packets, and the age of each such remaining active packet has increased by σ_t. Since there are no injections over successful streak $S(t)$, $\sigma_{t'} \geq \min_j\{s_j^t\} + \sigma_t > \sigma_t$. □

Lemma 17. *For an epoch beginning at time t, with constant probability, every streak is successful.*

Proof. Let t be any slot where $X(t) \geq 8$. Let $S(t_1), ..., S(t_k)$ be consecutive streaks such that k is the first index in these consecutive streaks where the contention drops below 8. Define $\tau_j = \sum_{i \leq j} \sigma_{t_i}$. Note that since $X(t) \geq 8$, then $\sigma_t \geq 2$. Then, by Lemma 13, the probability of a control failure over the interval $[t, t+\tau_k]$ is at most:

$$\left(\frac{1}{2}\right)^{b\cdot X(t)} + ... + \left(\frac{1}{2}\right)^{b\cdot X(t+\tau_{k-3})} + \left(\frac{1}{2}\right)^{b\cdot X(t+\tau_{k-2})} + \left(\frac{1}{2}\right)^{b\cdot X(t+\tau_{k-1})} + \left(\frac{1}{2}\right)^{b\cdot X(t+\tau_k)}$$

By assumption, $8 \leq X(t+\tau_{k-1})$ and, by Lemma 10, we know that $X(t+\tau_{k-1}) \leq \frac{3}{4}\cdot X(t+\tau_{k-2})$. Therefore $\frac{4}{3}\cdot 8 \leq \frac{4}{3}\cdot X(t+\tau_{k-1}) \leq X(t+\tau_{k-2})$. Similarly, we have that $X(t+\tau_{k-2}) \leq \frac{3}{4}\cdot X(t+\tau_{k-3})$. Therefore, $X(t+\tau_{k-3}) \geq \frac{4}{3}\cdot X(t+\tau_{k-2}) \geq (\frac{4}{3})^2 X(t+\tau_{k-1}) \geq (\frac{4}{3})^2\cdot 8$, and generally, $X(t+\tau_{k-j}) \geq (\frac{4}{3})^{j-1}\cdot 8$. Therefore, we can rewrite the terms as:

$$\begin{aligned}
&\leq \left(\frac{1}{2}\right)^{b\cdot X_t} + ... + \left(\frac{1}{2}\right)^{(\frac{4}{3})^2\cdot 8\cdot b} + \left(\frac{1}{2}\right)^{(\frac{4}{3})\cdot 8\cdot b} + \left(\frac{1}{2}\right)^{8\cdot b} + \left(\frac{1}{2}\right)^{b} \\
&\leq \sum_{j=0}^{\log_{4/3}(X_t)} \left(\frac{1}{2^{\Omega(b)}}\right)^{(4/3)^j} \\
&\leq \delta \text{ for any arbitrarily small constant } \delta > 0 \text{ depending only on sufficently large } b \text{ (depending only on } c)
\end{aligned}$$

Therefore, starting at time slot t, the probability that a control failure occurs in the interval $[t, t+\tau_k]$ defined by these consecutive streaks is at most δ. By Lemma 15, the probability of a restart is at most a constant $q/2$ depending only on d. Therefore, the epoch is successful with probability at least $1 - ((1-\frac{q}{2}) + \delta) = q/2 - \delta \geq \varepsilon$ for some constant $\varepsilon > 0$ depending only on c and d. □

Lemma 19. *In every iteration of the finite bad borrower game, the borrower borrows at least one dollar and at most n dollars.*

Proof. The fact that the borrower borrows at least one dollar follows by definition. Assume the borrower borrows n dollars, i.e., that the associated epoch lasts for at least n slots. Recall that the last streak in the

associated epoch must have been at least $n/2$ slots, and at the beginning of that final streak, the contention must have been at least 8 (or the epoch would have ended). Since the last streak is of length at least $n/2$, there must be a set of old packets with age $\geq n/2$ that collectively have contention at least 4 (by definition of a streak). This implies there must be at least $2n$ such packets, which is impossible, given the bound of n packets total. Thus, it is impossible to have an epoch of length n, and hence to borrow more than n dollars in an iteration of the finite bad borrower game. □

Lemma 20. *Over the entire execution of the bad borrower game, the expected number of dollars borrowed, in total, is $O(n)$.*

Proof. We analyze the dollars repaid in the following fashion: we assume that for every dollar lent, it is paid back with probability $p\alpha$. Notice, of course, that these random choices are correlated: for a given iteration, either an α fraction of the dollars are paid back (with probability p), or no dollars are paid back, with probability $1-p$. For a given iteration of the game, if there are k dollars borrowed, we see that the expected number of dollars repaid is $pk\alpha$, as expected.

We now ask, what is the expected number of dollars we have to lend in order for n dollars to be repaid? The answer is $n/(p\alpha)$, i.e., after $O(n)$ dollars have been borrowed, all n dollars have been repaid. In the last iteration, there can be at most n additional dollars lent (as part of the iteration where the last dollar is repaid), by Lemma 19, yielding an expected number of borrowed dollars of $O(n)+n$. □

Lemma 21. *For all iterations r of the infinite bad borrower game, there is some iteration $r' > r$ such that if the lender has lent k dollars through slot r', then the borrower has repaid at least $kp\alpha/2$ dollars, with probability 1.*

Proof. We can look at the random process as a one-dimensional biased random walk with variable step size. Let X_r be the value of the random walk in slot r, where $X_0 = 0$. Assume we lend x dollars in iteration r. With probability p, we succeed in slot r and hence we define $X_r = X_{r-1} - x + 2x/p$; otherwise, with probability $(1-p)$ we define $X_r = X_{r-1} - x$.

Notice that we have renormalized the random walk, so that every dollar paid back is worth $2/(p\alpha)$, i.e., if we have been paid back a $p\alpha/2$ fraction of the money, then our random walk is at zero. Thus if we show that the random walk is positive infinitely often, then we have completed the proof.

(Note that we cannot say that the random walk eventually remains always positive from some point on, as would be true of a simple constant-step-size random walk, because the adversary can always adjust the step size, for example, employing the following strategy: in each step where the random walk is positive, lend twice as much money until you lose and the random walk goes negative.)

To analyze this random walk, break the sequence of steps up into blocks in the following manner. If the previous block B_i ended with the random walk positive or zero, then block B_{i+1} contains only one step, i.e., the next step of the random walk. If the previous block B_i ended with the random walk negative, i.e., at $-b_i$, then continue the next block B_{i+1} up until the point where the adversary has cumulatively loaned $2b_i$; notice that in expectation, the random walk will increase by $2b_i$, ending the block positive at b_i. (Since the execution is infinite, and since the lender has to lend at least one dollar in each step, eventually every block will end.)

We now use a Hoeffding's inequality to show that with constant probability, the random walk returns to zero at the end of every block. First, if block B_{i+1} begins with the random walk zero or positive and takes only one step, then with constant probability that step is positive. Next, consider the case where B_{i+1} begins with the random walk negative, and lends at least $2b_i$ cumulatively throughout.

Let $Z_1, Z_2, \ldots, Z_k$ be the random variables associated with the change in value at each step of the random walk in the block, where in step Z_j the lender lends z_j dollars; with probability p, $Z_j = 2z_j/p - z_j$, and with probability $(1-p)$, $Z_j = -z_j$. Thus each Z_j has a bounded range $|Z_j|$ of size $2z_j/p$.

Let $\bar{z} = \sum_j z_j$, and recall that by construction, $\bar{z} \geq 2b_i$. Since $|Z_j| = 2z_j/p$, we know that $\sum |Z_j|^2 \leq 4\bar{z}^2/p^2$. Finally, we observe that for each Z_j, the expected value is z_j, and the expected value of the sum is $\bar{z}$.

Since the success or failure of epochs is independent (as packets are making independent choices in each slot), we can apply Hoeffding's inequality to lower bound the sum, where we choose $t = \bar{z}/2$:

$$\begin{aligned}\Pr((\sum_j Z_j - \bar{z}) \leq -t) &\leq e^{-\frac{2t^2}{\sum_j |Z_j|^2}} \\ &\leq e^{-2p^2t^2/(4\bar{z}^2)} \\ &\leq e^{-p^2/8}\end{aligned}$$

That is, with constant probability, the random walk gains at least $\bar{z}/2 \geq b_i$ during this block, and hence returns to zero.

To conclude the proof, we observe that since each block ends with the random walk returning to zero, there are an infinite number of points where the random walk returns to zero. (We cannot bound the length of time it takes to return to zero without first bounding the amount of money that can be lent.) Assume that at the end of block B_i, the random walk has returned to zero, and over the entire execution up until that point, the lender has lent k dollars. Since each dollar paid back causes the random walk to increase by $2/(p\alpha)$, this means that $kp\alpha/2$ dollars have been repaid, as required. □

Interstitial Slots

Lemma 23. *If an active packet u that becomes active in slot t_0 and resets in slot $t > t_0$. Then, all active packets injected in $[t_0, t]$ reset by slot t with probability 1.*

Proof. Let w be a packet injected in slot $t_0 \leq s < t$ that does not reset by t; we will derive a contradiction. Since u resets, a $7/8$-fraction of slots in $[t_0, s] \cup [s+1, t]$ are empty. We can divide these into two quantities of empty slots: $(\frac{7}{8} - \varepsilon)(t - t_0)$ and $\varepsilon(t - t_0)$ for $0 \leq \varepsilon \leq \frac{7}{8}$. Without loss of generality, assume there are $(\frac{7}{8} - \varepsilon)(t - t_0)$ empty slots in interval $[t_0, s]$ and there are $\varepsilon(t - s)$ empty slots in interval $[s + 1, t]$.

In the first interval, we must have $\frac{(7/8-\varepsilon)(t-t_0)}{s-t_0} < 7/8$ since otherwise u resets before t contradicting our setup. This implies that $(7/8 - \varepsilon)t < (7/8)s - \varepsilon\, t_0$. In the second interval, the fraction of empty slots is $\frac{\varepsilon(t-t_0)}{t-s-1} < 7/8$ where the inequality must hold since w does not reset. This implies that $(7/8-\varepsilon)t > (7/8)s-\varepsilon\, t_0+7/8$. Together, these inequalities imply that $(7/8)s-\varepsilon\, t_0+7/8 < (7/8-\varepsilon)t < (7/8)s-\varepsilon\, t_0$ and, therefore, we have a contradiction. □

Lemma 24. *With probability 1, the following hold:*

- *Any reset interval consists of at most a $\frac{7}{8}$-fraction of empty slots.*
- *Any interval of slots after a full reset consists of at most a $\frac{7}{8}$-fraction of empty slots.*
- *If there are no reset intervals, then the slots from when the first packet is broadcast and the last packet succeeds consists of at most a $\frac{7}{8}$-fraction of empty slots.*

Proof. Since we have removed from consideration all slots where the system size is zero, there is always at least one packet present at the first slot after a full reset. Let this packet be u from the above Lemma 23; this proves the first two claims. Similarly, if there are no reset intervals, let u be a packet present in the first slot. □

Lemma 25. *For a slot s, let $e_{\geq 2}$ denote the event where two or more packets are broadcast in s, and let $e_{=1}$ denote the event where one packet is broadcast in slot s. If s is an interstitial slot, then $Pr(e_{\geq 2}) = O(Pr(e_{=1}))$.*

Proof. The lemma follows immediately from Lemmas 7 and 8, since in interstitial slots, the contention is $O(1)$. □

Lemma 26. *There are at most $O(n)$ full, unjammed interstitial, slots in expectation.*

Proof. By the time that there are n full slots that have successful transmissions, the execution is over. And if we condition upon a given slot being full, there is a constant probability of a successful transmission by Lemma 25. Thus it takes $O(n)$ such slots, in expectation, before all n packets have successfully transmitted. □

Lemma 27. *If the adversary injects a total of n packets,* NOISEOFF *achieves expected constant throughput.*

Proof. We will argue that the expected number of slots for all the packets to finish is: $\mathrm{E}[T] = O(n + J)$ slots. We then observe that the expected throughput is $\mathrm{E}[\lambda] = \mathrm{E}[(n + J)/T]$, which by Jensen's inequality is $\geq O(1)$.

Let n_e denote the number of full slots over all unjammed epochs[8], and let n_i denote the number of full, unjammed interstitial slots. Let J denote the total number of jammed slots.

The execution can be divided up into reset intervals, or there exist no reset intervals; in either case, Lemma 24 implies that at most a $\frac{7}{8}$-fraction of the slots are empty over the execution. Therefore, over the entire execution, at most a $\frac{7}{8}$-fraction of the slots are empty. This implies that, over the execution, the number of empty interstitial slots is $O(n_e + n_i + J)$.

By Lemma 20, the finite bad borrower game implies that the number of epoch slots in unjammed epochs required to complete all n packets is $O(n)$ in expectation, therefore, $E[n_e] = O(n)$. As for the interstitial slots, Lemma 26 shows that $E[n_i] = O(n)$. Therefore, among unjammed epochs and unjammed interstitial slots, we conclude that the expected number of slots required for all n packets to succeed is $O(n)$.

Finally, we count the number of jammed slots. Since a jammed epoch is one in which at least $1/4$ of the slots are jammed, there are clearly at most $O(J)$ jammed epochs slots. Similarly, there are at most $O(J)$ jammed, non-empty insterstitial slots. There are also at most $O(J)$ slots in which the control channel is jammed (which can cause wasted time on the data channel if there are no active packets). Thus, there are at most $O(J)$ such slots otherwise unaccounted for. Thus we conclude that there are, in expectation, $O(n)$ unjammed epochs and unjammed interstitial slots, at most $O(J)$ jammed slots and jammed epochs, and $O(n + J)$ empty slots. □

Lemma 30. *In an infinite execution,* NOISEOFF *achieves constant throughput.*

Proof. For some slot t, let i_t^b be the number of successful broadcasts in unjammed interstitial slots prior to time t, and let i_t^c be the number of collisions in unjammed interstitial slots prior to time t.

We first argue that, with probability 1, from some point t onwards, for all $t' > t$: $i_t^c \leq O(i_t^b)$. Conditioned on the fact that there is at least one broadcast in an unjammed unterstitial round, let p be the probability of a successful broadcast and $q = 1 - p$ be the probability of a collision. We know from Lemma 25 that $q = O(p)$.

Define the following random walk: with probability q take a step to the left of size 1, and with probability p take a step to the right of size $2q/p$. Since $p \cdot (2q/p) > q$, by Corollary 29 we know that from some point on, this random walk is always positive.

Let t be a time slot that is after the last point where the random walk crosses the origin. We can then conclude that $i_t^c < i_t^b \cdot (2q/p)$. Since $2q/p = O(1)$, we conclude that $i_t^c = O(i_t^b)$.

Finally, we analyze the throughput. Fix any time t. Let $\hat{t}$ be the smallest time after t where the random walk defined above is positive. According to Lemma 21, there is a time $t' > \hat{t}$ where we have achieved constant throughput during the unjammed epochs, i.e., a constant fraction of the slots in unjammed epoch are broadcasts. By the analysis of the random walk, we conclude that a constant fraction of the non-empty,

[8] Recall that an epoch is unjammed if $< 1/4$ of its slots are unjammed.

unjammed interstitial rounds are broadcasts. By assumption, at most $O(J)$ slots are part of jammed epochs, and there are at most J jammed interstitial slots. There are at most J jammed control slots (which may cause delays on the data channel if there are no active packets). Finally, by Lemma 24, at most $7/8$ of the data channel slots in the entire execution are empty. Putting these pieces together yields constant throughput overall. □

Energy

Lemma 31. *With constant probability, a packet succeeds before it resets.*

Proof. In this proof, we grant the adversary even more power than given by the model—in each slot, the adversary is allowed to specify whether the slot is "covered", meaning that it is either jammed or some other packet transmits. The only thing the adversary does not control is the packet in question.

Starting from the time the packet becomes active, we divide time into windows $W_0, W_1, W_2, W_3, \ldots$, where window W_i has length 2^i. Note that if the first slot is covered, the packet cannot possibly reset until time 8 or later, which more than subsumes window W_1. Similarly, if at least 1 slot is also covered in window W_1, then the packet cannot possibly reset until after window W_2. In general, if at least half the slots are covered in each of the windows $W_0, W_1, \ldots, W_{i-1}$, then the packet either stays alive through W_i, or it succeeds sometime before the end of W_i—it cannot reset. Our argument thus proceeds inductively over windows, stopping at the first window W_i that is not at least half covered. In window W_i, the packet transmits independently in each data slot with probability at least $d/2^{i+2}$, where d is a constant specified in the protocol. Thus, if at least 2^{i-1} of the slots in W_i are left uncovered, the packet has either succeeded earlier, or it succeeds in window W_i with probability at least $1 - (1 - d/2^{i+2})^{2^{i-1}} \geq 1 - 1/e^{d/8}$, which is constant. □

Corollary 32. *For any positive integer k, the probability that a packet resets k times is at most $1/e^{\Theta(k)}$.*

Proof. The only observation we need is that for a particular packet, each of its lifetimes are nonoverlapping. Thus, each trial of Lemma 31 is independent. Each reset occurs with probability at most $1/e^{d/8}$, and hence the probability of k resets is at most $1/e^{kd/8} = 1/e^{\Theta(k)}$. □

Lemma 33. *In the first t slots following a packet's arrival, it expends $O(\log^2(t))$ energy in expectation.*

Proof. Consider any lifetime of the packet. The expected energy during a slot is equal to the packet's transmission probability, and hence the expected total energy is the sum of probabilities across all slots by linearity of expectation. The expected energy in a lifetime is thus at most $\sum_{s=1}^{t} \Theta(\ln s)/s = \Theta(\log^2 t)$, with the $\ln s$ arising from the high transmission probability in control slots.

We now compute the expected energy of the packet by using linearity of expectation across all lifetimes of the packet. In particular, the energy during the kth lifetime is 0 if the packet does not reset $k-1$ times, and hence applying Corollary 32 the expected energy of the kth lifetime is $1/e^{\Theta(k)} \cdot O(\log^2 t)$. Using linearity of expectation across all lifetimes, we get a total expected energy of $O(\log^2 t) \sum_{k=0}^{\infty} 1/e^{\Theta(k)} = O(\log^2 t)$. □

Lemma 34. *Consider the finite case, let n be the number of injected packets, and let J be the number of jammed slots. Then the expected energy each packet expends is $O(\log^2(n + J))$.*

Proof. Suppose the execution completes in time t. Then applying Lemma 33 yields an expected energy of $O(\log^2 t)$, as the packet must complete before the execution completes. Let T be the expected time of completion. From Markov's inequality, $t \geq T^i$ with probability at most $1/T^{i-1}$. Summing across all i, the expected energy becomes $O(\sum_{i=1}^{\infty}(1/T^i)\log^2(T^i)) = O(\log^2 T) \cdot \sum_{i=1}^{\infty}(i/T^i) = O(\log^2 T)$. Substituting in the expected makespan $T = O(n + J)$ from Lemma 27 concludes the theorem. □

Lemma 35. *Consider any time t in the infinite case at which we have λ throughput, for constant λ. Let J_t be the total number of jammed slots before t, and let η_t be the maximum contention prior to time t. Then the expected average energy per packet is $O(\log^2(\eta_t) + \log^2(J_t))$.*

Proof. The analysis is split into two cases: either $J_t \geq \lambda t/2$, or $J_t < \lambda t/2$. The first case is easy—Lemma 33 states that the expected energy per packet is $O(\log^2 t) = O(\log^2 J_t)$.

Suppose for the remainder that $J < \lambda t/2$. Then we divide the energy analysis here into three parts: (A) the jammed slots, (B) the unjammed slots for "young" packets, and (C) the unjammed slots for "old" packets. In parts A and B, we shall show that the expected energy per packet is at most $O(\log^2 J_t + \log^2 \eta_t)$, regardless of whether we have constant throughput. It is only in part C that we leverage the assumption that time t is a time at which we have constant throughput.

A. Consider a single lifetime of a specific packet. Our goal is to bound the energy expended by this packet during all J_t jammed slots. The energy is maximized if all J_t slots occur as early as possible in the packet's lifetime, in which case the expected energy is at most $O(\log^2 J_t)$ per lifetime. As in Lemma 33, we then apply Corollary 32 and linearity of expectation to get an expected energy of at most $O(\log^2 J_t) \sum_{k=0}^{\infty} 1/e^{\Theta(k)} = O(\log^2 J_t)$.

B. Consider a specific packet. We say that the packet is young during the first η_t^2 steps of its lifetime, during which it expends $O(\log^2 \eta_t^2) = O(\log^2 \eta_t)$ energy in expectation (Lemma 33). Summing across all lifetimes as above, we conclude that the contribution for young packets is $O(\log^2 \eta_t)$.

C. If a packet is not young, i.e., if its age is at least η_t^2, we say that it is old. Unlike parts A and B which analyze on a per-packet basis, this part analyzes the energy in aggregate. For every non-jammed slot, there are at most η_t packets in the system, and hence at most η_t packets are old. Each old packet transmits with probability at most $O(\ln(\eta_t^2)/\eta_t^2) = O(1/\eta_t)$, and hence the expected energy across all old packets in any slot is $O(1)$. Using linearity of expectation across all t slots, we have a total expected energy by old packets of $O(t)$.

Since we have constant throughput at time t, we know that at least λt slots are either jammed or successful transmissions. There are at most $\lambda t/2$ jammed slots by assumption, and hence there are at least $\lambda t/2 = \Omega(t)$ successful transmissions. We charge the $O(t)$ energy from old packets to these $\Omega(t)$ completed packets, for an additional $O(1)$ energy per successful packet. □